



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation and Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein together with the Statement of inventorship and of right to grant of a Patent (Form 7/77), which was subsequently filed.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed


Dated

1.8.02



1/77

Patents Act 1977
Rule 16



26JUN02 E728577-2 D00611
P01/7700 0.00-0214670.2

Request for grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920020011GB1		
2.	Patent application number (The Patent Office will fill in this part)	0214670.2		26 JUN 2002
3.	Full name, address and postcode of the or of each applicant (underline all surnames)	INTERNATIONAL BUSINESS MACHINES CORPORATION Armonk New York 10504 United States of America 519637001		
	Patents ADP number (if you know it)			
	If the applicant is a corporate body, give the country/state of its incorporation	State of New York United States of America		
4.	Title of the invention	ARRANGEMENT AND METHOD FOR UPDATE OF CONFIGURATION CACHE DATA		
5.	Name of your agent (if you have one)	R J Burt		
	"Address for Service" in the United Kingdom to which all correspondence should be sent (including the postcode)	IBM United Kingdom Limited Intellectual Property Department Hursley Park Winchester Hampshire SO21 2JN		
	Patents ADP number (if you know it)	919006		
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number	Country	Priority App No (if you know it)	Date of filing (day/month/year)
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date or the earlier application	No of earlier application	Date of filing (day/month/year)	



8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.)
- Yes

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description	6
Claim(s)	2
Abstract	1
Drawing(s)	2

22

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 7

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application


P J Stretton

Signature
R J Burt

25 June 2002
Date

12. Name and daytime telephone number of person to contact in the United Kingdom
- P J Stretton
01962 815830

The
Patent
Office




7/77

Patents Act 1977
Rule 15

Statement of inventorship and of right to grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920020011GB1	
2.	Patent application number (if you know it)	0214670.2	26 JUN 2002
3.	Full name of the or of each applicant	INTERNATIONAL BUSINESS MACHINES CORPORATION	
4.	Title of invention	ARRANGEMENT AND METHOD FOR UPDATE OF CONFIGURATION CACHE DATA	
5.	State how the applicant(s) derived the right from the inventor(s) to be granted a patent	By employment and by agreement	
6.	How many, if any, additional Patents Forms 7/77 are attached to this form?	1	
7.	I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.		
		 Signature R J Burt	25 June 2002 Date
8.	Name and daytime telephone number of person to contact in the United Kingdom	P J Stretton Tel: 01962 815830	



Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

David John CARR
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

Michael John JONES
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

Andrew KEY
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

REMINDER

Have you signed the form?

Enter full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Robert Bruce NICHOLSON
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

William James SCALES
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

Barry Douglas WHYTE
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

REMINDER

Have you signed the form?

ARRANGEMENT AND METHOD FOR UPDATE OF CONFIGURATION CACHE DATAField of the Invention

5 This invention relates to data storage systems (and subsystems) and methods for operation thereof.

Background of the Invention

10 In the field of this invention it is known that enhanced function storage subsystems (e.g., for a system of one or more disk drives) contain many logical resources. These resources can be managed by a storage adapter or controller. In this context, a logical resource is an entity which exposes a disk-like behaviour but which may have properties not
15 ordinarily associated with individual disk drives. These properties include enhanced performance, reliability, availability and data transformation operations such as flash copy. There may be a complex many-to-many relationship between logical resources and other logical resources or physical resources. For example a RAID-5 (Redundant Array of
20 Independent Disks) arrangement contains usually between 4 and 16 component disks. Data is striped across the components and a parity check is calculated and stored. The attributes of, and relationships between, these logical resources must be gathered and manipulated by software products to allow users to maintain their subsystem. Since the primary purpose of such
25 a storage subsystem is to perform Input/Output (IO) operations from operating system to disk drive(s) - IO has priority. This may be at the expense of user transactions that are sent to the subsystem in order to generate the details of the logical resources available in the subsystem. This in turn can result in very slow user response times - especially when
30 the subsystem is under heavy IO load.

 It is known to provide a cache of subsystem configuration data. This provides a centralised point of reference for user information relating to the subsystem. It is known from U.S. patent no. 5,895,493 that such a
35 system will gather information from various systems and provide a point for the user to access and obtain information about a group of systems.

 However, this approach has the disadvantage(s) that the system of the referenced patent does not describe how this data is updated. It relies on
40 the subsystems notifying the main cache of changes and updates and passing the data.

A need therefore exists for an arrangement and method of updating storage subsystem cache data wherein the abovementioned disadvantage(s) may be alleviated.

5 Statement of Invention

10 In accordance with a first aspect of the present invention there is provided an arrangement for update of cache data in a storage system, the arrangement comprising: memory means for holding data; cache means for
15 holding data associated with the memory means; means for preparing change data for updating the cache means; means for submitting a request for change to the memory means; means for receiving a signal from the memory means representative of completion of the request for change; and means for updating the cache means with the change data in response to the signal
15 indicating successful completion of the request for change.

Preferably, the data comprises configuration data.

20 Preferably, the storage system comprises a disk storage subsystem.

Preferably, the memory means is comprised in a disk adapter.

Preferably, the memory means is comprised in a disk controller.

25 The arrangement preferably further comprises interconnect means coupled between the memory means and the cache means for communicating the request for change, the signal representative of completion of the request for change.

30 Preferably, the interconnect means is also arranged to communicate transaction data.

Preferably, the the interconnect means comprises a device driver.

35 According to a second aspect of the present invention, there is provided a method for update of cache data in a storage system, the method comprising: providing memory means holding data; providing cache means holding data associated with the memory means; preparing change data for updating the cache means; submitting a request for change to the memory
40 means; receiving a signal from the memory means representative of completion of the request for change; and updating the cache means with the change data in response to the signal indicating successful completion of the request for change.

Preferably, the data comprises configuration data.

Preferably, the storage system comprises a disk storage subsystem.

5 Preferably, the memory means is provided in a disk adapter.

Preferably, the memory means is provided in a disk controller.

10 The method preferably further comprises providing interconnect means coupled between the memory means and the cache means for communicating the request for change, the signal representative of completion of the request for change.

15 Preferably, the interconnect means also communicates transaction data.

Preferably, the interconnect means comprises a device driver.

20 In a third aspect, the present invention provides a computer program element comprising computer program means for performing the method of the second aspect, and preferably for performing the additional preferred features of the method of the second aspect.

Brief Description of the Drawings

25

One arrangement and method of updating configuration cache data incorporating the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

30

FIG. 1 is a schematic diagram showing possible routes for user requests in a disk storage subsystem having a host service routine or agent containing a cache and accessing the storage subsystem via an interconnect;

35

FIG. 2 shows a flow chart illustrating a write-through cache method with two phase commit used in the system of FIG. 1.

Description of Preferred Embodiment

40

Enhanced function storage subsystems (e.g., for a system of one or more disk drives) contain many logical resources. Each logical resource generally relates to a particular function, for example a RAID-5 logical resource, a flash copy logical resource, etc. The attributes associated with these logical resources must be gathered by software products to allow users to maintain their subsystem. Since the primary purpose of such a

storage subsystem is to perform Input/Output (IO) operations from operating system to disk drive(s) - IO has priority. This may be at the expence of user transactions that are sent to the subsystem in order to generate the details of the logical resources available in the subsystem. This in turn
5 can result in very slow user response times - especially when the subsystem is under heavy IO.

In summary, the preferred embodiment of this invention provides a way in which the configuration software running in the host can cache the results of the logical resource configuration discovery transactions and hence re-use the same transaction data when applicable. The present invention is based on the realisation by the inventor that in general
10 subsystem logical resources do not change, and that therefore the cache can be used in preference to "talking" to the subsystem to find out what it has in terms of logical resources. This results in greatly reduced user response times, and reduces the non-essential workload on the subsystem. In particular the preferred embodiment of this invention provides a method of cache update involving a 2-phase commit of the data.
15

The configuration cache can be implemented in various ways, the preferred embodiment being illustrated with respect to FIG. 1.
20

The cache 110 is part of an agent or service 120 that runs in the host that connects to the storage subsystem. This service, when
25 initialised, can request all the logical resource information from the storage subsystem 130 via an interconnect 140 which may be the same interconnect that is used for I/O transactions or may be another interconnect specifically provided for the purposes of configuration. The logical configuration data that is returned is saved in the cache for subsequent use. The storage subsystem also returns a Configuration Version
30 Number (CVN). This number is the current CVN stored (130C) by the storage subsystem, and hence 'version stamps' the current configuration cache. The storage subsystem increments this number whenever it detects a change in its logical configuration. In this way the software can ask the storage
35 subsystem its CVN, check this against the CVN (110C) that was stored in the cache and act accordingly.

It will be understood that, although as shown in FIG. 1 communication between the service and the subsystem occurs via the interconnect 140, this
40 interconnect may in practice be provided in the form of a device driver in an adapter-based solution where the device driver communicates with an adapter (not shown) of the storage subsystem. Alternatively, the interconnect may communicate with a controller (not shown) of the storage subsystem. It will further be understood that the arrangement shown in

FIG. 1 represents both "in band" communication where the service uses the same route as I/O transactions and "out of band" communication over perhaps Ethernet or some other interconnect.

For example, where a single host contains an adapter and the required "service" software is running in the host:

- a) The service requests the entire logical configuration from the adapter.
- b) The service receives the transaction data (stores it in the cache) and a CVN = 123.
- c) Some time later the service receives a request from one of its clients that involves generating a "view" of a logical resource.
- d) The service requests the adapter CVN (130C)
- e) The adapter returns a CVN = 123
- f) The service checks this against its cache (110C), they are the same and so returns the cached "view" to the client. The User Request therefore follows 'Route B' (as shown by dotted line in FIG. 1) to produce the Response.

Some time later the service receives a new request for another "view" of a logical resource:

- g) The service requests the adapter CVN (130C)
- h) The adapter returns a CVN = 125
- i) The service checks this against its cache (110C), they are different and so the service then submits the relevant transactions to the adapter so as to rebuild the "view" of the logical resource.
- j) The adapter returns the transaction data. The User Request therefore follows 'Route A' (as shown by dashed line in FIG. 1) to produce the Response.

At this stage, the cache is now possibly incoherent, and in the idle loop the contents can be updated.

It will be understood that the system as described above in relation to FIG. 1 is well known and need not be described in further detail. It will also be understood that the cache invalidation technique required to cope with possible incoherence of the cache is well known and also need not be described further.

Referring now also to FIG. 2, a write-through cache with two phase commit technique is used for logical configuration cache update.

When the user submits a configuration change transaction, the details of the request are constructed and are built into a request in the form that the subsystem requires (step 210). At this time a delta (change data) is prepared for the cache and is temporarily stored as a change pending for the logical resource cache (step 220) - this step constitutes phase 1 of the two-phase commit procedure.

The request is then sent to the subsystem for processing (step 230).

Upon receipt of the completed request the return code for the request is checked (step 240) (the return code may conveniently be an integer value that indicates either success or failure of a particular request - in the case of a failure the return code may be decoded into a failure reason):

- k) if the request completed successfully the delta is applied to the cache and the cache is now up to date with the contents of the subsystem (step 260) - this constitutes phase 2 of the two-phase commit procedure;
- l) if the request failed, the delta is destroyed (step 260) because there is no need for it to be applied to the cache.

The transaction is then completed as normal (step 270).

It will be understood that the method of update of configuration cache data described above provides the advantage that known changes to the subsystem do not require an invalidate/rebuild style operation on the cache. This is especially important in the cache outlined above, where a change will invalidate the entire cache.

It will be appreciated that the method described above for update of configuration cache data will typically be carried out in software running on a processor (not shown), and that the software may be provided as a computer program element carried on any suitable data carrier (also not shown) such as a magnetic or optical computer disc.

It will also be appreciated that various modifications to the embodiment described above will be apparent to a person of ordinary skill in the art.

CLAIMS

1. An arrangement for update of cache data in a storage system, the arrangement comprising:

memory means for holding data;

cache means for holding data associated with the memory means;

means for preparing change data for updating the cache means;

means for submitting a request for change to the memory means;

means for receiving a signal from the memory means representative of completion of the request for change; and

means for updating the cache means with the change data in response to the signal indicating successful completion of the request for change.

2. The arrangement of claim 1 wherein the data comprises configuration data.

3. The arrangement of claim 1 or 2 wherein the storage system comprises a disk storage subsystem.

4. The arrangement of claim 3 wherein the memory means is comprised in a disk adapter.

5. The arrangement of claim 3 wherein the memory means is comprised in a disk controller.

6. The arrangement of any preceding claim further comprising interconnect means coupled between the memory means and the cache means for communicating the request for change, the signal representative of completion of the request for change.

7. The arrangement of claim 4 wherein the interconnect means is also arranged to communicate transaction data.

8. The arrangement of claim 6 or 7 wherein the interconnect means comprises a device driver.

9. A method for update of cache data in a storage system, the method comprising:

providing memory means holding data;

providing cache means holding data associated with the memory means;

5 preparing change data for updating the cache means;

submitting a request for change to the memory means; receiving a
signal from the memory means representative of completion of the request
for change; and

10 updating the cache means with the change data in response to the
signal indicating successful completion of the request for change.

10. The method of claim 9 wherein the data comprises configuration data.

15 11. The method of claim 9 or 10 wherein the storage system comprises a
disk storage subsystem.

20 12. The method of claim 13 wherein the memory means is provided in a disk
adapter.

13. The method of claim 12 wherein the memory means is provided in a disk
controller.

25 14. The method of any one of claims 9-13 further comprising providing
interconnect means coupled between the memory means and the cache means for
communicating the request for change, the signal representative of
completion of the request for change.

30 15. The method of claim 14 wherein the interconnect means also
communicates transaction data.

16. The method of claim 14 or 15 wherein the interconnect means comprises
a device driver.

35 17. A computer program element comprising computer program means for
performing the method of any one of claims 9-16.

40 18. An arrangement, for update of cache data in a storage system,
substantially as hereinbefore described with reference to the accompanying
drawings.

19. A method, for update of cache data in a storage system, substantially
as hereinbefore described with reference to the accompanying drawings.

ABSTRACT**ARRANGEMENT AND METHOD FOR UPDATE OF CONFIGURATION CACHE DATA**

5 An arrangement and method for update of configuration cache data in a
disk storage subsystem in which a cache memory (110) is updated using
two-phase (220, 250) commit technique.

10 This provides the advantage that known changes to the subsystem do
not require an invalidate/rebuild style operation on the cache. This is
especially important where a change will invalidate the entire cache.

FIG. 1

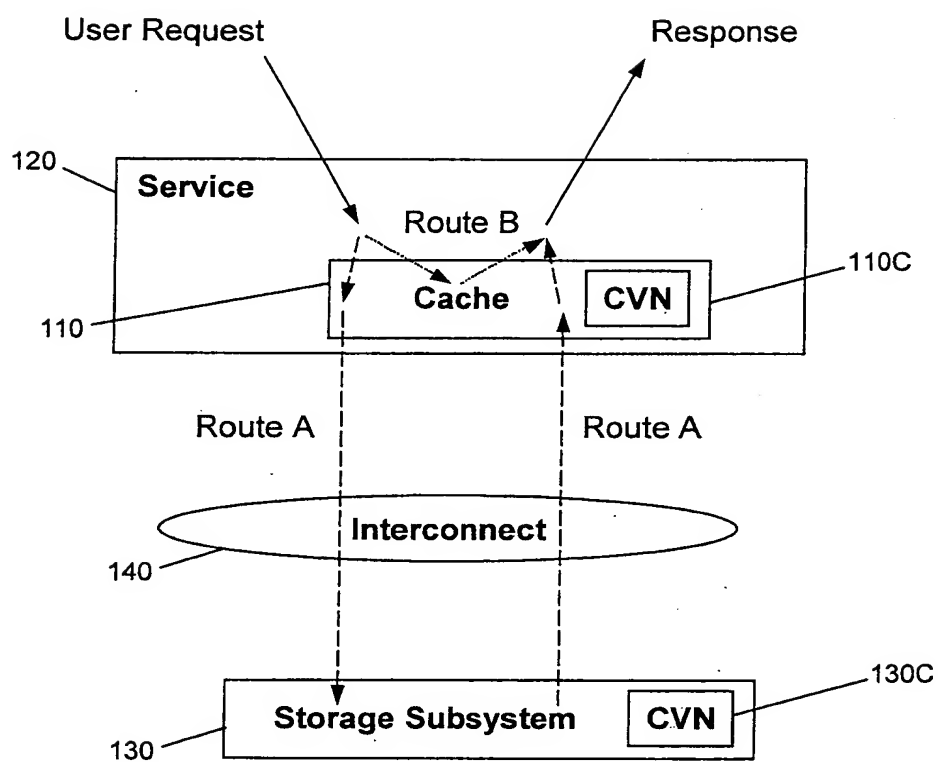


FIG. 2

